

Towards Explainable RAG - Interpreting the Influence of Retrieved Passages on Generation: Resource Efficiency and Performance Trade-offs

Wyatt Holland, Alec Hawkins, Brett Keller | Review Article | Published 2026-08-27

ABSTRACT

Efficiency claims should state which resources are saved, what performance is exchanged, and whether the trade-off remains acceptable at operational scale. This structured evidence review evaluates "Towards Explainable RAG: Interpreting the Influence of Retrieved Passages on Generation" alongside nine author-disjoint, topically matched publications in reliable retrieval-augmented generation. It compares construct definitions, evaluation choices, operating assumptions, and reported limitations instead of treating bibliographic similarity as empirical equivalence. Viewed through resource efficiency and performance trade-offs, the map separates claims supported by the available record from questions that still require full-text extraction, replication, or new experiments. The synthesis is interpretive rather than meta-analytic and therefore does not present a pooled effect estimate or a new causal result. The resulting agenda reports compute, memory, energy, latency, and maintenance costs beside task performance at realistic scale.

KEYWORDS

reliable retrieval-augmented generation; resource efficiency and performance trade-offs; evidence synthesis; reproducibility; research evaluation

Research Context

Sang (2025) [1] provides the focal point for this review. Its topic is consequential because progress in reliable retrieval-augmented generation depends not only on a proposed method, material, dataset, or workflow, but also on the credibility of the evidence chain that connects design decisions to reported outcomes. The present article therefore asks a deliberately bounded question: how should the focal work be interpreted when the primary lens is resource efficiency and performance trade-offs? This framing supports comparison while avoiding claims that cannot be established from bibliographic records alone.

The review follows a structured evidence-mapping protocol. First, the focal work defines the problem boundary. Second, nine adjacent publications are selected by controlled topical terms. Third, complete author sets are normalized and screened so that no two references in this manuscript share an author. Fourth, each item is assigned an explicit analytical role. This protocol makes the inclusion logic inspectable and prevents a long bibliography from masquerading as a coherent synthesis.

Evidence Map

Evidence node 2 is Hassan et al. (2026) [2], which addresses "Agentic Self-RAG: Multi-Agent Reasoning for Self-Correcting Retrieval-Augmented Generation". Its controlled title terms place it directly within reliable retrieval-augmented generation; in this review it is used to test how the focal argument behaves when viewed through resource efficiency and performance trade-offs.

Evidence node 3 is Fasolino (2026) [3], which addresses "In RAG We Trust? Measuring Robustness of Retrieval-Augmented Generation Under Document Poisoning". Its controlled title terms place it directly within reliable retrieval-augmented generation; in this review it is used to test how the focal argument behaves when viewed through resource efficiency and performance trade-offs.

Evidence node 4 is Samal (2026) [4], which addresses "A Theoretical Analysis of Self-Contained Retrieval-Augmented Generation with Small Language Models". Its controlled title terms place it directly within reliable retrieval-augmented generation; in this review it is used to test how the focal argument behaves when viewed through resource efficiency and performance trade-offs.

Evidence node 5 is Pi (2024) [5], which addresses "Efficient Information Retrieval and Response Generation with Retrieval-Augmented Generation (RAG)". Its controlled title terms place it directly within reliable retrieval-augmented generation; in this review it is used to test how the focal argument behaves when viewed through resource efficiency and performance trade-offs.

Evidence node 6 is Xue et al. (2025) [6], which addresses "A Comparative Study of Retrieval-Augmented Generation, Graph Retrieval-Augmented Generation, and Fine-Tuned Large Language Models for Fire Engineering Knowledge Retrieval". Its controlled title terms place it directly within reliable retrieval-augmented generation; in this review it is used to test how the focal argument behaves when viewed through resource efficiency and performance trade-offs.

Evidence node 7 is Bose (2025) [7], which addresses "Introduction to Retrieval-Augmented Generation (RAG)". Its controlled title terms place it directly within reliable retrieval-augmented generation; in this review it is used to test how the focal argument behaves when viewed through resource efficiency and performance trade-offs.

Evidence node 8 is Zhai (2025) [8], which addresses "SAM-RAG: An Self-adaptive Framework for Multimodal Retrieval-Augmented Generation". Its controlled title terms place it directly within reliable retrieval-augmented generation; in this review it is used to test how the focal argument behaves when viewed through resource efficiency and performance trade-offs.

Evidence node 9 is Gudipati et al. (2026) [9], which addresses "CityCopilot-X: A Real-Time Explainability Panel for Retrieval-Augmented Generation (RAG)". Its controlled title terms place it directly within reliable retrieval-augmented generation; in this review it is used to test how the focal argument behaves when viewed through resource efficiency and performance trade-offs.

Evidence node 10 is Kau (2024) [10], which addresses "Understanding Retrieval Pitfalls: Challenges Faced by Retrieval Augmented Generation (RAG) models". Its controlled title terms place it directly within reliable retrieval-augmented generation; in this review it is used to test how the focal argument behaves when viewed through resource efficiency and performance trade-offs.

Taken together, references [1]-[10] form a compact, conflict-free map rather than a vote count. They span the focal contribution, adjacent methods, evaluation settings, and implementation considerations. Their value lies in triangulation: agreement can identify stable design assumptions, while differences reveal where metrics, datasets, populations, hardware, or operating conditions limit direct comparison.

Analytical Framework

The first analytical layer is construct alignment. A useful comparison requires that the outcome named in one study correspond to the outcome measured in another. For manuscript 01-P06-029, this means distinguishing nominally similar terms from operationally equivalent measures. The second layer is evidence strength. Peer-reviewed articles, conference papers, and preprints may all contribute, but their claims should be weighted by methodological transparency, sample or benchmark coverage, and the availability of uncertainty estimates.

The third layer concerns transfer. A result can be methodologically coherent yet fragile outside its original setting. Under the lens of resource efficiency and performance trade-offs, transfer should be tested along at least four axes: data composition, task definition, resource envelope, and decision cost. The fourth layer is failure analysis. Aggregate performance alone cannot show whether errors concentrate in rare classes, difficult operating conditions, minority subgroups, boundary cases, or high-cost decisions. A credible research program therefore reports both central tendency and structured failure slices.

The fifth layer is reproducibility. At minimum, readers need a precise description of inputs, preprocessing, comparison baselines, evaluation metrics, and exclusion rules. Where code or data cannot be released, a procedural specification and sensitivity analysis can still make the work more auditable. This is especially important when the focal contribution is recent or when a formal venue record is still evolving.

Implications for Research and Practice

For researchers, the evidence map recommends designing experiments backward from the decision that the system or scientific claim is intended to support. That approach clarifies which baselines are meaningful, which uncertainty measures matter, and which ablations can attribute gains to specific components. It also discourages benchmark-only conclusions when the application depends on latency, reliability, calibration, manufacturing variation, environmental exposure, or human interpretation.

For practitioners, the key implication is staged adoption. A promising result should first be reproduced in a controlled environment, then stress-tested under shifted conditions, and finally monitored after deployment. Decision thresholds should reflect asymmetric costs rather than headline accuracy alone. Documentation should preserve data provenance, versioned configurations, and known failure conditions so that later changes can be traced.

Limitations and Research Agenda

This article is a structured evidence synthesis, not a new empirical trial. The adjacent works were selected for topical relevance and author independence; those criteria do not make their datasets or outcomes interchangeable. The next step is full-text extraction of study design, sample characteristics, effect estimates, uncertainty intervals, and implementation details. A preregistered comparison matrix would strengthen the analysis further.

Three questions follow. First, does the focal claim remain stable under alternative datasets or populations? Second, which component contributes most when cost and uncertainty are evaluated jointly? Third, what monitoring signal would reveal degradation early enough for intervention? Answering these questions would turn the present evidence map into a cumulative research program centered on resource efficiency and performance trade-offs.

References

1. Sang, Y. (2025). Towards Explainable RAG: Interpreting the Influence of Retrieved Passages on Generation. 2025 4th International Conference on Robotics, Artificial Intelligence and Intelligent Control (RAIIC), 397-400. <https://doi.org/10.1109/raaic65850.2025.11170170>
2. Hassan, S.-B., Abdullah., & Abbas, M. (2026). Agentic Self-RAG: Multi-Agent Reasoning for Self-Correcting Retrieval-Augmented Generation. 2026 International Conference on IT and Industrial Technologies (ICIT), 1-6. <https://doi.org/10.1109/icit68548.2026.11577708>
3. Fasolino, I. (2026). In RAG We Trust? Measuring Robustness of Retrieval-Augmented Generation Under Document Poisoning. . <https://doi.org/10.32388/7hxxqe>
4. Samal, M.-P. (2026). A Theoretical Analysis of Self-Contained Retrieval-Augmented Generation with Small Language Models. . <https://doi.org/10.36227/techrxiv.177219989.96070478/v1>
5. Pi, W. (2024). Efficient Information Retrieval and Response Generation with Retrieval-Augmented Generation (RAG). . <https://doi.org/10.59350/q2pq3-0fv85>
6. Xue, X., Zhang, G., Jiang, L., & Liu, C. (2025). A Comparative Study of Retrieval-Augmented Generation, Graph Retrieval-Augmented Generation, and Fine-Tuned Large Language Models for Fire Engineering Knowledge Retrieval. . <https://doi.org/10.2139/ssrn.5316632>
7. Bose, R. (2025). Introduction to Retrieval-Augmented Generation (RAG). Mastering Retrieval-Augmented Generation, 3-32. <https://doi.org/10.1007/979-8-8688-1808-01>
8. Zhai, W. (2025). SAM-RAG: An Self-adaptive Framework for Multimodal Retrieval-Augmented Generation. 2025 International Joint Conference on Neural Networks (IJCNN), 1-8. <https://doi.org/10.1109/ijcnn64981.2025.11227819>
9. Gudipati, S.-K., Mishra, N.-A., Ankam, G., Akkiseti, S.-R., Mohammed, A., & Ponugoti, S. (2026). CityCopilot-X: A Real-Time Explainability Panel for Retrieval-Augmented Generation (RAG). 2026 IEEE 5th International Conference on AI in Cybersecurity (ICAIC), 1-9. <https://doi.org/10.1109/icaic67076.2026.11395791>
10. Kau, A. (2024). Understanding Retrieval Pitfalls: Challenges Faced by Retrieval Augmented Generation (RAG) models. . <https://doi.org/10.59350/xcq3s-jvk04>